AFRL-AFOSR-VA-TR-2015-0315

Apprenticeship Learning for Robotic Control

Pieter Abbeel
REGENTS OF THE UNIVERSITY OF CALIFORNIA THE

10/09/2015
Final Report

DISTRIBUTION A: Distribution approved for public release.

Air Force Research Laboratory
AF Office Of Scientific Research (AFOSR)/ RTA2
Arlington, Virginia 22203
Air Force Materiel Command

REPORT DOCUMENTATION PAGE	Form Approved OMB No. 0704-0188

1. REPORT DATE *(DD-MM-YYYY)*	**2. REPORT TYPE**	**3. DATES COVERED** *(From - To)*
9/26/2015	Final Report	2012/7/1 - 2015/6/30

4. TITLE AND SUBTITLE

Apprenticeship Learning for Robotic Control

5a. CONTRACT NUMBER

5b. GRANT NUMBER

FA9550-12-1-0345

5c. PROGRAM ELEMENT NUMBER

6. AUTHOR(S)

Pieter Abbeel

5d. PROJECT NUMBER

5e. TASK NUMBER

5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)

REGENTS OF THE UNIVERISTY OF CALIFORNIA
THE UNIVERSITY OF CALIFORNIA, BERKELEY
2150 SHATTUCK AVE, RM 313
BERKELEY, CA 94704-5940

8. PERFORMING ORGANIZATION REPORT NUMBER

033376-003

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)

OFV OF NAVAL RSCH (ONRRO) SAN DIEGO
SAN DIEGO REGIONAL OFFICE
140 SYLVESTER RD BLDG 140 RM 218
SAN DIEGO, CA 92106-3521

10. SPONSOR/MONITOR'S ACRONYM(S)

AFOSR

11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT

Public Distribution

13. SUPPLEMENTARY NOTES

14. ABSTRACT

Our research had three main thrusts: (i) Optimization-based motion planning: Traditional approaches to motion planning tend to slow down in high-dimensional spaces and for curvature constrained problems. We developed an optimization-based approach that can efficiently solve such problems. (ii) Belief space planning: To account for uncertainty about the environment and robot belief space planning tries to find a plan that optimizes (on expectation) the sequence of probability distributions that would result from that plan. As a consequence belief space planning anticipates (and accounts for) sensory information, the expected magnitude of perturbations, and controllability. We made contributions in handling collisions, handling occlusions, scaling up belief space planning, going beyond unimodal beliefs, and hierarchical belief space planning. (iii) Learning from demonstrations: We developed an approach to generalize demonstrated motion in training scenarios to test scenarios. At the core of our approach is non-rigid registration to warp the training scene onto the test scene. While registration is only concerned with the objects and their environment, we show that it is possible to meaningfully extrapolate the registration and to also warp the robot gripper trajectories from

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:			**17. LIMITATION OF ABSTRACT**	**18. NUMBER OF PAGES**	**19a. NAME OF RESPONSIBLE PERSON**
a. REPORT	**b. ABSTRACT**	**c. THIS PAGE**			Pieter Abbeel
U	U	U	UU		**19b. TELEPHONE NUMBER** *(Include area code)* 650 387 8115

Reset

DISTRIBUTION A: Distribution approved for public release

INSTRUCTIONS FOR COMPLETING SF 298

1. REPORT DATE. Full publication date, including day, month, if available. Must cite at least the year and be Year 2000 compliant, e.g. 30-06-1998; xx-06-1998; xx-xx-1998.

2. REPORT TYPE. State the type of report, such as final, technical, interim, memorandum, master's thesis, progress, quarterly, research, special, group study, etc.

3. DATES COVERED. Indicate the time during which the work was performed and the report was written, e.g., Jun 1997 - Jun 1998; 1-10 Jun 1996; May - Nov 1998; Nov 1998.

4. TITLE. Enter title and subtitle with volume number and part number, if applicable. On classified documents, enter the title classification in parentheses.

5a. CONTRACT NUMBER. Enter all contract numbers as they appear in the report, e.g. F33615-86-C-5169.

5b. GRANT NUMBER. Enter all grant numbers as they appear in the report, e.g. AFOSR-82-1234.

5c. PROGRAM ELEMENT NUMBER. Enter all program element numbers as they appear in the report, e.g. 61101A.

5d. PROJECT NUMBER. Enter all project numbers as they appear in the report, e.g. 1F665702D1257; ILIR.

5e. TASK NUMBER. Enter all task numbers as they appear in the report, e.g. 05; RF0330201; T4112.

5f. WORK UNIT NUMBER. Enter all work unit numbers as they appear in the report, e.g. 001; AFAPL30480105.

6. AUTHOR(S). Enter name(s) of person(s) responsible for writing the report, performing the research, or credited with the content of the report. The form of entry is the last name, first name, middle initial, and additional qualifiers separated by commas, e.g. Smith, Richard, J, Jr.

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES). Self-explanatory.

8. PERFORMING ORGANIZATION REPORT NUMBER. Enter all unique alphanumeric report numbers assigned by the performing organization, e.g. BRL-1234; AFWL-TR-85-4017-Vol-21-PT-2.

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES). Enter the name and address of the organization(s) financially responsible for and monitoring the work.

10. SPONSOR/MONITOR'S ACRONYM(S). Enter, if available, e.g. BRL, ARDEC, NADC.

11. SPONSOR/MONITOR'S REPORT NUMBER(S). Enter report number as assigned by the sponsoring/ monitoring agency, if available, e.g. BRL-TR-829; -215.

12. DISTRIBUTION/AVAILABILITY STATEMENT. Use agency-mandated availability statements to indicate the public availability or distribution limitations of the report. If additional limitations/ restrictions or special markings are indicated, follow agency authorization procedures, e.g. RD/FRD, PROPIN, ITAR, etc. Include copyright information.

13. SUPPLEMENTARY NOTES. Enter information not included elsewhere such as: prepared in cooperation with; translation of; report supersedes; old edition number, etc.

14. ABSTRACT. A brief (approximately 200 words) factual summary of the most significant information.

15. SUBJECT TERMS. Key words or phrases identifying major concepts in the report.

16. SECURITY CLASSIFICATION. Enter security classification in accordance with security classification regulations, e.g. U, C, S, etc. If this form contains classified information, stamp classification level on the top and bottom of this page.

17. LIMITATION OF ABSTRACT. This block must be completed to assign a distribution limitation to the abstract. Enter UU (Unclassified Unlimited) or SAR (Same as Report). An entry in this block is necessary if the abstract is to be limited.

Report Type: Final Report

Primary Contact Email: pabbeel@cs.berkeley.edu

Primary Contact Phone Number: (650) 387 8115

Organization / Institution name: UC Berkeley

Contract/Grant Title: (YIP-12) Apprenticeship Learning for Robotic Control
Contract/Grant #: FA9550-12-1-0345

PI Name: Pieter Abbeel

Program Manager: James Lawton

Reporting Period: 2012/7/1 – 2015/6/30

Abstract:

Our research had three main thrusts: (i) Optimization-based motion planning: Traditional approaches to motion planning tend to slow down in high-dimensional spaces and for curvature constrained problems. We developed an optimization-based approach that can efficiently solve such problems. (ii) Belief space planning: To account for uncertainty about the environment and robot belief space planning tries to find a plan that optimizes (on expectation) the sequence of probability distributions that would result from that plan. As a consequence belief space planning anticipates (and accounts for) sensory information, the expected magnitude of perturbations, and controllability. We made contributions in handling collisions, handling occlusions, scaling up belief space planning, going beyond unimodal beliefs, and hierarchical belief space planning. (iii) Learning from demonstrations: We developed an approach to generalize demonstrated motion in training scenarios to test scenarios. At the core of our approach is non-rigid registration to warp the training scene onto the test scene. While registration is only concerned with the objects and their environment, we show that it is possible to meaningfully extrapolate the registration and to also warp the robot gripper trajectories from training scene to test scene. Our approach is particularly appealing for the manipulation of deformable objects, which present the robot with a very high-dimensional state space and large amounts of variability, making them particularly challenging for robots to manipulate. It has enabled autonomous knot tying for a wide range of knot-types and automation of some simplified suturing tasks.

Annual Accomplishments:

Our research had three main thrusts: (i) Optimization-based motion planning, (ii) Belief space planning, (iii) Learning from demonstrations.

(i) Optimization-based motion planning

Traditional approaches to motion planning tend to slow down in high-dimensional spaces and for curvature constrained problems (such as needle steering and ribbon routing). We developed an optimization-based approach that can efficiently solve such problems [58, 68, 74]. In this work we introduced a new optimization-based approach for robotic motion planning among obstacles. Like CHOMP (Covariant Hamiltonian Optimization for Motion Planning), our algorithm can be used to find collision-free trajectories from naïve, straight-line initializations that might be in collision. At the core of our approach are (a) a sequential convex optimization procedure, which penalizes collisions with a hinge loss and increases the penalty coefficients in an outer loop as necessary, and (b) an efficient formulation of the no-collisions constraint that directly considers continuous-time safety. Our algorithm is implemented in a software package called TrajOpt. We validated our approach through a series of experiments comparing TrajOpt with CHOMP and randomized planners from OMPL, with regard to planning time and path quality. We considered motion planning for 7 DOF robot arms, 18 DOF full-body robots, statically stable walking motion for the 34 DOF Atlas humanoid robot, and physical experiments with the 18 DOF PR2. We also applied TrajOpt to plan curvature-constrained steerable needle trajectories in the SE(3) configuration space and multiple non-intersecting curved channels within 3D-printed implants for intracavitary brachytherapy.

(ii) Belief space planning

Uncertainty about the environment is unavoidable in unstructured settings. Uncertainty about the robot itself can also be a factor, especially as robots become cheaper and cheaper. Belief space planning algorithms try to find a plan that optimizes (on expectation) the sequence of probability distributions that would result from that plan. As a consequence belief space planning anticipates (and accounts for) sensory information, the expected magnitude of perturbations, and controllability.

Building on advances in sequential convex optimization our most recent work is able to find (locally) optimal paths in Gaussian belief spaces. Our key contributions have been in handling probabilistic collision checking for articulated robots [62], in handling sensors with a limited field of view and occlusions [69], and in exploiting the structure of Gaussian belief space planning problems to enable scaling up to significantly higher dimensional problems (i.e., higher number of degrees of

freedom of the robot, higher number of objects, and higher number of landmarks) than was possible before [75].

Gaussian belief space planning, while surprisingly effective, nevertheless has significant limitations as it is only able to reason about unimodal beliefs. In [C89] we have started to investigate multimodal belief representations in the context of grasping occluded objects.

Recent work has made headway on integration of task and motion planning in state space. We developed Interfaced Belief Space Planning (IBSP): a modular approach to task and motion planning in belief space. We use a task-independent interface layer to combine an off-the-shelf classical planner with motion planning and inference. We determinize the problem under the maximum likelihood observation assumption to obtain a deterministic representation where successful plans generate goal-directed observations. We leverage properties of maximum likelihood observation determinizations to obtain a simple representation of (optimistic) belief space dynamics that is well suited to planning. Our interface is implemented with standard belief state queries, requiring only the ability to sample, compute unnormalized likelihoods, and compute maximum likelihood states. Our contribution is a novel algorithm for task and motion planning in belief space that has minimal dependence on the details of the inference engine used. IBSP can work with a broad class of black box state estimators, with zero changes to the algorithm. We validate our approach in simulated tasks for the PR2 that account for continuous state, different types of initial state distributions, and negative observations.

(iii) Learning from demonstrations

In [C64] we have started to develop the beginnings of a theory that allows for generalization from training scenarios to test scenarios. At the core of our approach is non-rigid registration to warp the training scene (where the demonstration happens) onto the test scene (which is a previously unseen situation where the robot has to autonomously succeed). While registration is only concerned with the objects and their environment, we show that it is possible to meaningfully extrapolate the registration and to also warp the robot gripper trajectories from training scene to test scene. Our approach is particularly appealing for the manipulation of deformable objects, which present the robot with a very high-dimensional state space and large amounts of variability, making them particularly challenging for robots to manipulate. Our approach has enabled autonomous knot tying for a wide range of knot-types and starting configurations. It has enabled automation of some simplified suturing tasks. Figure 1 illustrates some of our experiments.

Figure 1: Autonomous knot-tying, suturing, and folding of a t-shirt. Key insight in this project was how to build on existing techniques for non-rigid registration to enable generalization of motions by warping them appropriately from training scene to test scene.

Our approach as discussed thus far only considers robot gripper pose trajectories, but often what matters forces and torques the robot applies. For instance, tightening a knot requires pulling the ends, flattening an article of clothing requires smoothing out wrinkles, and erasing a whiteboard requires applying downward pressure. In [C92] we show how the variation between multiple demonstrations of variants of the same task can be used to determine when, and to which extent, forces than well pose matter. This results in a learned variable impedance control strategy that trades off force and position errors, providing the right level of compliance for each motion.

Improved registration [C87]: The first step in our approach is a non-rigid registration between training scene and test scene. The current registration only considers a geometric representation of train and test scenes and it considers all parts of the scenes equally important. In [C87] we show how to incorporate surface normals. Future work will study how to incorporate visual (in addition to geometric) features.

Archival Publications (published) during Reporting Period:

Numbering matches number on my webpage: www.cs.berkeley.edu/~pabbeel

[106] Modular Task and Motion Planning in Belief Space,
Dylan Hadfield-Menell, Edward Groshev, Rohan Chitnis, Pieter Abbeel.
In the proceedings of the *28th IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Hamburg, Germany, September 2015. (pdf)

[92] Learning Force-Based Manipulation of Deformable Objects from Multiple Demonstrations,
Alex Lee, Henry Lu, Abhishek Gupta, Sergey Levine, Pieter Abbeel.

In the proceedings of the *IEEE International Conference on Robotics and Automation (ICRA)*, 2015. (pdf)

[87] A Non-Rigid Point and Normal Registration Algorithm with Applications to Learning from Demonstrations,
Alex Lee, Max Goldstein, Shane Barratt, Pieter Abbeel.
In the proceedings of the *IEEE International Conference on Robotics and Automation (ICRA)*, 2015. (pdf)

[89] Active Exploration using Trajectory Optimization for Robotic Grasping in the Presence of Occlusions,
Gregory Kahn, Peter Sujan, Sachin Patil, Shaunak D. Bopardikar, Julian Ryde, Ken Goldberg, Pieter Abbeel.
In the proceedings of the *IEEE International Conference on Robotics and Automation (ICRA)*, 2015. (pdf)

[75] Scaling up Gaussian Belief Space Planning through Covariance-Free Trajectory Optimization and Automatic Differentiation,
Sachin Patil, Greg Kahn, Michael Laskey, John Schulman, Ken Goldberg, Pieter Abbeel.
In the proceedings of the *11th International Workshop on the Algorithmic Foundations of Robotics (WAFR)*, Aug. 2014. (pdf)

[72] Motion Planning with Sequential Convex Optimization and Convex Collision Checking,
John Schulman, Yan Duan, Jonathan Ho, Alex Lee, Ibrahim Awwal, Henry Bradlow, Jia Pan, Sachin Patil, Ken Goldberg, Pieter Abbeel.
In the *International Journal of Robotics Research (IJRR)*, Vol. 33, No. 9, pp. 1251-1270, Aug. 2014. (pdf)

[69] Gaussian Belief Space Planning with Discontinuities in Sensing Domains,
Sachin Patil, Yan Duan, John Schulman, Ken Goldberg, Pieter Abbeel.
In the proceedings of the *IEEE International Conference on Robotics and Automation (ICRA)*, 2014. (pdf)

[68] Planning Locally Optimal, Curvature-Constrained Trajectories in 3D using Sequential Convex Optimization,
Yan Duan, Sachin Patil, John Schulman, Ken Goldberg, Pieter Abbeel.

In the proceedings of the *IEEE International Conference on Robotics and Automation (ICRA)*, 2014. (pdf)

[64] Learning from Demonstrations through the Use of Non-Rigid Registration,
John Schulman, Jonathan Ho, Cameron Lee, Pieter Abbeel.
In the proceedings of the *16th International Symposium on Robotics Research (ISRR)*, Dec 2013. (pdf)

[63] A Case Study of Trajectory Transfer through Non-Rigid Registration for a Simplified Suturing Scenario,
John Schulman, Ankush Gupta, Sibi Venkatesan, Mallory Tayson-Frederick, Pieter Abbeel.
In the proceedings of the *26th IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Nov 2013. (pdf, talk video, supplementary materials)

[62] Sigma Hulls for Gaussian Belief Space Planning for Imprecise Articulated Robots amid Obstacles,
Alex Lee, Yan (Rocky) Duan, Sachin Patil, John Schulman, Zoe McCarthy, Jur van den Berg, Ken Goldberg, Pieter Abbeel.
In the proceedings of the *26th IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Nov 2013. (pdf, talk video)

[58] Finding Locally Optimal, Collision-Free Trajectories with Sequential Convex Optimization,
John D. Schulman, Jonathan Ho, Alex Lee, Ibrahim Awwal, Henry Bradlow and Pieter Abbeel.
In the proceedings of *Robotics: Science and Systems (RSS)*, Jul 2013. (pdf, videos, code)

Changes in Research Objective, if any:

None

Change in AFOSR program manager, if any:

Dr. Jay Myung left as of 2012/5/31. When starting off the AFOSR was working to bring in a new program officer to manage Dr. Jay Myung's portfolio. In the interim, Dr. Robert Bonneau (robert.bonneau@afosr.af.mil; 703-696-9545), with the help of Ms. Stephanie Bruce (stephanie.bruce@afosr.af.mil; 703-588-0664), oversaw the portfolio. Ultimately Dr. James Lawton became the new program officer, and still currently is.

Extensions Granted or Milestones Slipped, if any:

None

Include any new discoveries, inventions, or patent disclosures during this reporting period (if none, report none)

None

1.

1. Report Type

Final Report

Primary Contact E-mail
Contact email if there is a problem with the report.

pabbeel@cs.berkeley.edu

Primary Contact Phone Number
Contact phone number if there is a problem with the report

6503878115

Organization / Institution name

UC Berkeley

Grant/Contract Title
The full title of the funded effort.

Apprenticeship Learning for Robotic Control

Grant/Contract Number
AFOSR assigned control number. It must begin with "FA9550" or "F49620" or "FA2386".

FA9550-12-1-0345

Principal Investigator Name
The full name of the principal investigator on the grant or contract.

Pieter Abbeel

Program Manager
The AFOSR Program Manager currently assigned to the award

James Lawton

Reporting Period Start Date

07/01/2012

Reporting Period End Date

06/30/2015

Abstract

Our research had three main thrusts: (i) Optimization-based motion planning: Traditional approaches to motion planning tend to slow down in high-dimensional spaces and for curvature constrained problems. We developed an optimization- based approach that can efficiently solve such problems. (ii) Belief space planning: To account for uncertainty about the environment and robot belief space planning tries to find a plan that optimizes (on expectation) the sequence of probability distributions that would result from that plan. As a consequence belief space planning anticipates (and accounts for) sensory information, the expected magnitude of perturbations, and controllability. We made contributions in handling collisions, handling occlusions, scaling up belief space planning, going beyond unimodal beliefs, and hierarchical belief space planning. (iii) Learning from demonstrations: We developed an approach to generalize demonstrated motion in training scenarios to test scenarios. At the core of our approach is non-rigid registration to warp the training scene onto the test scene. While registration is only concerned with the objects and their environment, we show that it is possible to meaningfully extrapolate the registration and to also warp the robot gripper trajectories from training scene to test scene. Our approach is particularly appealing for the manipulation of deformable objects, which present the robot with a very high-dimensional state space and large amounts of variability, making them particularly challenging for robots to manipulate. It has enabled autonomous knot tying for a wide range of knot-types and automation of some simplified

suturing tasks.

Distribution Statement

This is block 12 on the SF298 form.

Distribution A - Approved for Public Release

Explanation for Distribution Statement

If this is not approved for public release, please provide a short explanation. E.g., contains proprietary information.

SF298 Form

Please attach your SF298 form. A blank SF298 can be found here. Please do not password protect or secure the PDF The maximum file size for an SF298 is 50MB.

AFD-070820-035.pdf

Upload the Report Document. File must be a PDF. Please do not password protect or secure the PDF . The maximum file size for the Report Document is 50MB.

AFOSR-YIP_final_report_2015_09_30.pdf

Upload a Report Document, if any. The maximum file size for the Report Document is 50MB.

Archival Publications (published) during reporting period:

Numbering matches number on my webpage: www.cs.berkeley.edu/~pabbeel

[106] Modular Task and Motion Planning in Belief Space,
Dylan Hadfield-Menell, Edward Groshev, Rohan Chitnis, Pieter Abbeel.
In the proceedings of the 28th IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Hamburg, Germany, September 2015.

[92] Learning Force-Based Manipulation of Deformable Objects from Multiple Demonstrations,
Alex Lee, Henry Lu, Abhishek Gupta, Sergey Levine, Pieter Abbeel.
In the proceedings of the IEEE International Conference on Robotics and Automation (ICRA), 2015.

[87] A Non-Rigid Point and Normal Registration Algorithm with Applications to Learning from Demonstrations,
Alex Lee, Max Goldstein, Shane Barratt, Pieter Abbeel.
In the proceedings of the IEEE International Conference on Robotics and Automation (ICRA), 2015.

[89] Active Exploration using Trajectory Optimization for Robotic Grasping in the Presence of Occlusions,
Gregory Kahn, Peter Sujan, Sachin Patil, Shaunak D. Bopardikar, Julian Ryde, Ken Goldberg, Pieter Abbeel.
In the proceedings of the IEEE International Conference on Robotics and Automation (ICRA), 2015.

[75] Scaling up Gaussian Belief Space Planning through Covariance-Free Trajectory Optimization and Automatic Differentiation,
Sachin Patil, Greg Kahn, Michael Laskey, John Schulman, Ken Goldberg, Pieter Abbeel.
In the proceedings of the 11th International Workshop on the Algorithmic Foundations of Robotics (WAFR), Aug. 2014.

[72] Motion Planning with Sequential Convex Optimization and Convex Collision Checking,
John Schulman, Yan Duan, Jonathan Ho, Alex Lee, Ibrahim Awwal, Henry Bradlow, Jia Pan, Sachin Patil, Ken Goldberg, Pieter Abbeel.
In the International Journal of Robotics Research (IJRR), Vol. 33, No. 9, pp. 1251-1270, Aug. 2014.

[69] Gaussian Belief Space Planning with Discontinuities in Sensing Domains,
Sachin Patil, Yan Duan, John Schulman, Ken Goldberg, Pieter Abbeel.
In the proceedings of the IEEE International Conference on Robotics and Automation (ICRA), 2014.

[68] Planning Locally Optimal, Curvature-Constrained Trajectories in 3D using Sequential Convex Optimization,
Yan Duan, Sachin Patil, John Schulman, Ken Goldberg, Pieter Abbeel.
In the proceedings of the IEEE International Conference on Robotics and Automation (ICRA), 2014.

[64] Learning from Demonstrations through the Use of Non-Rigid Registration,
John Schulman, Jonathan Ho, Cameron Lee, Pieter Abbeel.
In the proceedings of the 16th International Symposium on Robotics Research (ISRR), Dec 2013.

[63] A Case Study of Trajectory Transfer through Non-Rigid Registration for a Simplified Suturing Scenario,
John Schulman, Ankush Gupta, Sibi Venkatesan, Mallory Tayson-Frederick, Pieter Abbeel.
In the proceedings of the 26th IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Nov 2013.

[62] Sigma Hulls for Gaussian Belief Space Planning for Imprecise Articulated Robots amid Obstacles,
Alex Lee, Yan (Rocky) Duan, Sachin Patil, John Schulman, Zoe McCarthy, Jur van den Berg, Ken Goldberg, Pieter Abbeel.
In the proceedings of the 26th IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Nov 2013.

[58] Finding Locally Optimal, Collision-Free Trajectories with Sequential Convex Optimization,
John D. Schulman, Jonathan Ho, Alex Lee, Ibrahim Awwal, Henry Bradlow and Pieter Abbeel.
In the proceedings of Robotics: Science and Systems (RSS), Jul 2013.

Changes in research objectives (if any):

n/a

Change in AFOSR Program Manager, if any:

Dr. Jay Myung left as of 2012/5/31. When starting off the AFOSR was working to bring in a new program officer to manage Dr. Jay Myung's portfolio. In the interim, Dr. Robert Bonneau (robert.bonneau@afosr.af.mil; 703-696-9545), with the help of Ms. Stephanie Bruce (stephanie.bruce@afosr.af.mil; 703-588-0664), oversaw the portfolio. Ultimately Dr. James Lawton became the new program officer, and still currently is.

Extensions granted or milestones slipped, if any:

n/a

AFOSR LRIR Number

LRIR Title

Reporting Period

Laboratory Task Manager

Program Officer

Research Objectives

Technical Summary

Funding Summary by Cost Category (by FY, $K)

	Starting FY	FY+1	FY+2
Salary			
Equipment/Facilities			
Supplies			
Total			

Report Document

Report Document - Text Analysis

Report Document - Text Analysis

Appendix Documents

2. Thank You

E-mail user

Sep 29, 2015 18:56:30 Success: Email Sent to: pabbeel@cs.berkeley.edu